

Connectivity Preserving Network Transformers*

Othon Michail and Paul G. Spirakis

Abstract The Population Protocol model is a distributed model that concerns systems of very weak computational entities that cannot control the way they interact. The model of Network Constructors is a variant of Population Protocols capable of (algorithmically) constructing abstract networks. Both models are characterized by a *fundamental inability to terminate*. In this work, we investigate the minimal strengthenings of the latter model that could overcome this inability. Our main conclusion is that *initial connectivity of the communication topology* combined with the ability of the protocol to *transform the communication topology* and the ability of a node to *detect when its degree is equal to a small constant*, plus either a *unique leader* or the ability of *detecting common neighbors*, are sufficient to guarantee not only *termination* but also the *maximum computational power that one can hope for in this family of models*. In particular, the model, under these minimal assumptions, *computes with termination any symmetric predicate computable by a Turing Machine of space $\Theta(n^2)$* .

Othon Michail

Department of Computer Science, University of Liverpool, Liverpool, UK &
Computer Technology Institute and Press “Diophantus” (CTI), Patras, Greece
e-mail: Othon.Michail@liverpool.ac.uk

Paul G. Spirakis

Department of Computer Science, University of Liverpool, Liverpool, UK &
Computer Technology Institute and Press “Diophantus” (CTI), Patras, Greece
e-mail: P.Spirakis@liverpool.ac.uk

* Supported in part by (i) the project “Foundations of Dynamic Distributed Computing Systems” (FOCUS) which is implemented under the “ARISTEIA” Action of the Operational Programme “Education and Lifelong Learning” and is co-funded by the European Union (European Social Fund) and Greek National Resources and (ii) the FET EU IP project MULTIPLEX under contract no 317532. The full paper on which this chapter is based, can be found at <http://arxiv.org/abs/1512.02832>.

1 Introduction

A dynamic distributed computing system is a system composed of distributed computational processes in which the structure of the communication network between the processes changes over time. In one extreme, the processes cannot control and cannot accurately predict the modifications of the communication topology. Typical such examples are mobile distributed systems in which the mobility is *external* to the processes and is usually provided by the environment in which the system operates. For example, it could be a system of cell phones following the movement of the individuals carrying them or a system of nanosensors flowing in the human circulatory system. This type of mobility is known as *passive* (see e.g. [AAD⁺06]). On the other extreme, dynamicity may be a sole outcome of the algorithm executed by the processes. Typical examples are systems in which the processes are equipped with some *internal* mobility mechanism, like mobile robotic systems and, in general, any system with the ability to algorithmically modify the communication topology. This type of mobility is known as *active* mobility (see e.g. [WCG⁺13] for active self-assembly, [SY99, DFSY15, CKLWL09] for mobile robots, and [ACD⁺11] for reconfigurable (nano)robotics under physical constraints). Recently, there is an interest in *intermediate* (or *hybrid*) systems. One such type, consists of systems in which the processes are passively mobile but still they are equipped with an internal active mechanism that allows them to have a partial (algorithmic) control of the system's dynamicity.

The intermediate model that guides our study here, is the *network constructors* model introduced in [MS14]. In this model, there are n extremely weak processes, computationally equivalent to anonymous finite automata, that usually have very limited knowledge of the system (e.g. they do not know its size). The processes move passively and interact in pairs whenever two of them come sufficiently close to each other. This part of the system's dynamicity is not controlled and cannot be (completely) predicted by the processes and is modeled by assuming an adversary scheduler that in every step selects a pair of processes to interact. The adversary is typically restricted to be *fair* so that it cannot forever block the system's progress (e.g. by keeping two parts of the system forever disconnected). Fairness is sufficient for analyzing the correctness of protocols for specific tasks. If additionally an estimate of the running time is desired, a typical assumption is that the scheduler is a uniform random one (which is fair with probability 1 [CDF⁺09] and also corresponds to the dynamicity patterns of well-mixed solutions). But in this model, there is also an internal source of dynamicity. In particular, the processes can algorithmically connect and disconnect to each other during their pairwise interactions. This can be viewed either as a physical bonding mechanism, as e.g. in reconfigurable robotics and molecular (e.g. DNA) self-assembly, or as a virtual record of local connectivity, as e.g. in a social network where a participant keeps track of and can regularly update the set of his/her associates. This allows the processes to control the construction and maintenance of a network or a shape in an uncontrolled and unpredictable dynamic environment.

The network from which the scheduler picks interactions between processes and develops the uncontrolled interaction pattern is called the *interaction network*. At the same time, the processes, by connecting and disconnecting to each other, develop another network, the (*algorithmically*) *constructed network*, which is a subnetwork of the interaction network. In the most abstract setting, the interaction network is the clique K_n throughout the execution, no matter what the protocol does (e.g. no matter how the protocol modifies the constructed network). In this case, the scheduler can in every step (throughout the course of the protocol) pick any possible pair of processes to interact, independently of the constructed network.² This is precisely the setting of [MS14] and also the one that we will consider in the present work.³ But even if the interaction network is always a clique independently of the constructed network, the ability of the processes to construct a network may still allow them to counterbalance the adversary’s power. For example, if the processes manage somehow to self-organize into a spanning network G , then it might be possible for them to ignore all interactions that occur over the non-links of G and thus force the actual communication pattern to be consistent with the constructed network.

The existing literature on distributed network construction [MS14, Mic15] has almost absolutely focused on the setting in which all processes are initially disconnected and the goal is for them to algorithmically self-organize into a desired (usually spanning or of size at least some required function of n) stable network or shape. In [MS14], the authors presented simple and efficient direct constructors and lower bounds for several basic network construction problems such as spanning line, spanning ring, and spanning star and also generic constructors capable of constructing a large class of networks by simulating a Turing Machine (abbreviated “TM” throughout). One of the main results was that for every graph language L that is decidable by a $O(\sqrt{l})$ -space ($l + O(\sqrt{l})$, resp.) TM, where $l = \Theta(n^2)$ is the binary length of the input of the simulated TM, there is a protocol that constructs L equiprobably with useful space $\lfloor n/2 \rfloor$ ($\lfloor n/3 \rfloor$, resp.), where the *useful space* is defined as a lower bound on the order of the output network (the rest of the nodes being used as auxiliary and thrown away eventually as *waste*). In [Mic15], a geometrically constrained variant was studied, where the formed network and the allowable interactions must respect the structure of the 2-dimensional (or 3-dimensional) grid network. The main result was a *terminating* protocol counting the size n of the system with high probability (abbreviated “w.h.p.” throughout). This protocol was then used as a subroutine of universal constructors, establishing that the nodes can self-assemble w.h.p. into arbitrarily complex shapes while still being capable to terminate after completing the construction.

² A convenient way to think of this setting is to imagine a clique graph with its edges labeled from $\{0, 1\}$. Then, in this case, the clique is the interaction network while its subgraph induced by the edges labeled 1 is the constructed network.

³ On the other hand, it is possible, and plausible w.r.t. several application scenarios, that the set of available interactions at a given step actually depends on the constructed network. Such a case was considered in [Mic15], where the constructed network is always a subnetwork of the grid network and two processes can only interact if a connection between them would preserve this requirement. So, in that case, the set of available interactions is, in every step, constrained by the network that has been constructed by the protocol so far.

1.1 Our Approach and Contribution

The main goal of this work is to investigate minimal strengthenings of the population protocol and network constructors models that can maximize their computational power, also rendering them capable to terminate. To this end, we consider (for the first time in network constructors) the case in which the initial configuration of the edges is not the one in which all edges are *inactive* (i.e. those that are in state 0). In particular, we assume that the initial configuration of the edges can be any configuration in which the *active* (i.e. those that are in state 1) edges form a *connected graph spanning the set of processes*.⁴ The initial configuration of the nodes is either, as in [MS14], the one in which all nodes are initially in the same state, e.g. in an initial state q_0 , or (whenever needed) the one in which all nodes begin from q_0 apart from a pre-elected unique leader that begins from a distinct initial leader-state l_0 . This choice is motivated by the fact that without some sort of bounded initial disconnectivity we can only hope for global computations and constructions that are *eventually stabilizing* (and not *terminating*), because a component can guess neither the number of components not encountered yet nor an upper bound on the time needed to interact with another one of them ([MS15] overcomes this by assuming that the nodes know some upper bound on this time, while [Mic15] overcomes this by assuming a uniform random scheduler and a unique leader and by restricting correctness to be w.h.p.).

Next, observe that if the protocol is not allowed to modify the state of the edges, then the assumption of initial connectivity alone does not add any computational power to the model (in the worst case). For if we ignore for a while the ability of the model to modify the state of the edges, what we have is a model equivalent to classical population protocols [AAD⁺06] on a restricted interaction graph [AAC⁺05] (observe that the model can ignore the interactions that occur over inactive edges). Though there are some restricted interaction graphs, like the spanning line, that dramatically increase the computational power of the model (in this case making it equivalent to a TM of linear space), still there others, like the spanning star, on which the power of the model is as low as the power of classical population protocols on a clique interaction graph [CMNS13], which, in turn, is equal to the rather small class of *semilinear predicates* [AAER07]. As we have allowed any possible connected initial set of active edges, the spanning star inclusive, the initial configuration of the edges alone (without any edge modifications) is not sufficient for strengthening the model.

Our discussion so far, suggests to consider at the same time initial connectivity (or, more generally, bounded initial disconnectivity) and the ability of the protocol to modify the state of the edges, with the hope of increasing the computational power. Unfortunately, even with this additional assumption, non-trivial terminating computation is still impossible (this is proved in Proposition 3, in Section 3.3). An immediate way to appreciate this, is to notice that a clique does not provide more

⁴ Active and inactive *edges* are not to be confused with active and passive *mobility*. An edge is said to be *active* if its state is 1 and it is said to be *inactive* if its state is 0.

information than an empty network about the size of the system. Even worse, if a node’s initial active degree is unbounded (as e.g. is the case for the center of a spanning star), then it is not clear even whether the stabilizing constructors that assume initial disconnectivity (as in [MS14]) can be adapted to work. Actually, it could be the case, that *without additional assumptions* initial connectivity may even decrease the power of the model (we leave this as an interesting open problem). For example, it could be simpler to construct a spanning line if the initial active network is empty (i.e. all edges are inactive) than if it is a clique (i.e. all edges are active). Even if it would turn out that the model does not become any weaker, we still cannot avoid the aforementioned impossibility of termination and the maximum that we can hope for is an *eventually stabilizing* universal constructor, as the one of [MS14].

We now add to the picture a very minimal and natural, but extremely powerful, additional assumption that, combined with our assumptions so far, will lead us to a stronger model. In particular, we equip the nodes with the ability to detect some small local degrees. For a concrete example, assume that a node can detect when its active degree is equal to 0 (otherwise it only knows that its degree is at least 1). A first immediate gain, is that we can now directly simulate any constructor that assumes an empty initial network (e.g. the constructors of [MS14]): every node initially deactivates the active edges incident to it until its local active degree becomes for the first time 0, and only when this occurs the node starts participating in the simulation. So, even though a node does not know its initial degree (which is due to the fact that a node in this model is a finite automaton with a state whose size is independent of the size of the system), it can still detect when it becomes equal to 0. At that point, the node does not have any active edges incident to it, therefore it can start executing the constructor that assumes an empty initial network.

Our main finding in this work, is that the initial connectivity guarantee together with the ability to modify the network and to detect small local degrees (combined with either a pre-elected leader or a natural mechanism that allows two nodes to tell whether they have a neighbor in common), are sufficient to obtain the *maximum computational power* that one can hope for in this family of models. In particular, the resulting model can compute *with termination* any symmetric predicate⁵ computable by a *TM of space* $\Theta(n^2)$, and no more than this, i.e. it is an exact characterization. The symmetry restriction can only be dropped by UIDs or by any other means of knowing and maintaining an ordering of the nodes’ inputs. This power is maximal because the distributed space of the system is $\Theta(n^2)$, so we cannot hope for computations exploiting more space. The substantial improvement compared to [MS14, MCS11a] is that the universal computations are now *terminating* and not just *eventually stabilizing*. It is interesting to point out that the additional assumptions and mechanisms are minimal, in the sense that the removal of each one of them leads to either an impossibility of termination or to a substantial decrease in the computational power.

In Section 2, we discuss further related literature. Section 3 brings together all definitions and basic facts that are used throughout the chapter. In particular, in Sec-

⁵ Essentially, a predicate in this type of models is called *symmetric* (or *commutative*) if permuting the input symbols does not affect the predicate’s outcome.

tion 3.1 we formally define the model of network constructors under consideration, Section 3.2 formally defines the transformation problems that are considered in this work, and Section 3.3 provides some basic impossibility results and a lower bound on the time needed to transform any network to a spanning line. In Section 4, we study the case in which there is a pre-elected unique leader and give two protocols for the problem, the Online-Cycle-Elimination protocol and the time-optimal Line-Around-a-Star protocol. Then, in Section 5, we try to drop the unique leader assumption. First, in Section 5.1 we show that, without additional assumptions, dropping the unique leader leads to a strong impossibility result. In face of this negative result, in Section 5.2 we minimally strengthen the model with a common neighbor detection mechanism and give a correct terminating protocol. Finally, in Section 6 we conclude and give further research directions that are opened by our work.

2 Further Related Work

The model considered in this chapter belongs to the family of population protocol models. The population protocol model [AAD⁺06] was originally developed as a model of highly dynamic networks of simple sensor nodes that cannot control their mobility. The first papers focused on the computational capabilities of the model which have now been almost completely characterized. In particular, if the interaction network is complete, i.e. one in which every pair of processes may interact, then the computational power of the model is equal to the class of the *semilinear predicates* (and the same holds for several variations) [AAER07]. Semilinearity persists up to $o(\log \log n)$ local space but not more than this [CMN⁺11]. If additionally the connections between processes can hold a state from a finite domain (note that this is a stronger requirement than the active/inactive that the present work assumes) then the computational power dramatically increases to the commutative subclass of $\mathbf{NSPACE}(n^2)$ [MCS11a]. The latter constitutes the mediated population protocol (MPP) model, which was the first variant of population protocols to allow for states on the edges. For introductory texts to these models, the interested reader is encouraged to consult [AR09] and [MCS11b].

Based on the MPP model, [MS14] restricted attention to binary edge states and regarded them as a physical (or virtual, depending on the application) bonding mechanism. This gave rise to a “hybrid” self-assembly model, the network constructors model, in which the actual dynamicity is passive and due to the environment but still the protocol can construct a desired network by activating and deactivating appropriately the connections between the nodes. The present chapter essentially investigates the computational power of the network constructors model under the assumption that a connected spanning active topology is provided initially and also initiates the study of the *distributed network reconfiguration problem*. Recently, [Mic15] studied a geometrically constrained variant of network constructors in which the interaction network is not complete but rather it is constrained by the existing shapes (every shape that can be formed being a sub-network of the 2D or 3D

grid network). Interestingly, apart from being a model of computation, population protocols are also closely related to chemical systems. In particular, Doty [Dot14] has recently demonstrated their formal equivalence to *chemical reaction networks* (CRNs), which model chemistry in a *well-mixed solution*.

3 Preliminaries

3.1 The Model

The model under consideration is the network constructors model of [MS14] with the only essential difference being that in [MS14] the initial configuration was always (apart from the network replication problem) the one in which all edges are inactive, while in this work the initial configuration can be any configuration in which the active edges form a spanning connected network. Still, we give a detailed presentation of the model for self-containment.

Definition 1. A *Network Constructor* (NET) is a distributed protocol defined by a 4-tuple $(Q, q_0, Q_{out}, \delta)$, where Q is a finite set of *node-states*, $q_0 \in Q$ is the *initial node-state*, $Q_{out} \subseteq Q$ is the set of *output node-states*, and $\delta : Q \times Q \times \{0, 1\} \rightarrow Q \times Q \times \{0, 1\}$ is the *transition function*. When required, also a special *initial leader-state* $l_0 \in Q$ may be defined.

If $\delta(a, b, c) = (a', b', c')$, we call $(a, b, c) \rightarrow (a', b', c')$ a *transition* (or *rule*) and we define $\delta_1(a, b, c) = a'$, $\delta_2(a, b, c) = b'$, and $\delta_3(a, b, c) = c'$. A transition $(a, b, c) \rightarrow (a', b', c')$ is called *effective* if $x \neq x'$ for at least one $x \in \{a, b, c\}$ and *ineffective* otherwise. When we present the transition function of a protocol we only present the effective transitions. Additionally, we agree that the *size* of a protocol is the number of its states, i.e. $|Q|$.

The system consists of a population V_I of n distributed *processes* (also called *nodes* when clear from context). In the generic case, there is an underlying *interaction graph* $G_I = (V_I, E_I)$ specifying the permissible interactions between the nodes. Interactions in this model are always pairwise. In this work, unless otherwise stated, G_I is a *complete undirected interaction graph*, i.e. $E_I = \{uv : u, v \in V_I \text{ and } u \neq v\}$, where $uv = \{u, v\}$. When we say that all nodes in V_I are initially *identical*, we mean that all nodes begin from the initial node-state q_0 . In case we assume the existence of a unique leader, then there is a $u \in V_I$ beginning from the initial leader-state l_0 and all other $v \in V_I \setminus \{u\}$ begin from the initial node-state q_0 (which in this case may also be called the *initial nonleader-state*).

A central assumption of the model is that edges have binary states. An edge in state 0 is said to be *inactive* while an edge in state 1 is said to be *active*. In almost all problems studied in [MS14] (apart from the replication problem), all edges were initially inactive. Though we shall also consider this case in the present chapter, our main focus is on a different setting in which the protocol begins its execution on a

precomputed set of active edges provided by some adversary. Formally, there is an input set of edges $E \subseteq E_I$, such that all $e \in E$ are initially active and all $e' \in E_I \setminus E$ are initially inactive. The set E defines the *input graph* $G = (V_I, E)$, also called the *initial active topology/graph*. Throughout this work, unless otherwise stated, we assume that the initial active topology is *connected*, which means that the active edges form a connected graph spanning V_I . This is a restriction imposed on the adversary selecting the input. In particular, the adversary is allowed to choose any initial set of active edges E (in a worst-case manner), subject to the constraint that E defines a connected graph on the whole population.

Execution of the protocol proceeds in discrete steps. In every step, a pair of nodes uv from E_I is selected by an *adversary scheduler* and these nodes interact and update their states and the state of the edge joining them according to the transition function δ . In particular, we assume that, for all distinct node-states $a, b \in Q$ and for all edge-states $c \in \{0, 1\}$, δ specifies either (a, b, c) or (b, a, c) . So, if a, b , and c are the states of nodes u, v , and edge uv , respectively, then the unique rule corresponding to these states, let it be $(a, b, c) \rightarrow (a', b', c')$, is applied, the edge that was in state c updates its state to c' and if $a \neq b$, then u updates its state to a' and v updates its state to b' , if $a = b$ and $a' = b'$, then both nodes update their states to a' , and if $a = b$ and $a' \neq b'$, then the node that gets a' is drawn equiprobably from the two interacting nodes and the other node gets b' .

A *configuration* is a mapping $C : V_I \cup E_I \rightarrow Q \cup \{0, 1\}$ specifying the state of each node and each edge of the interaction graph. Let C and C' be configurations, and let u, v be distinct nodes. We say that C goes to C' via encounter $e = uv$, denoted $C \xrightarrow{e} C'$, if $(C'(u), C'(v), C'(e)) = \delta(C(u), C(v), C(e))$ or $(C'(v), C'(u), C'(e)) = \delta(C(v), C(u), C(e))$ and $C'(z) = C(z)$, for all $z \in (V_I \setminus \{u, v\}) \cup (E_I \setminus \{e\})$. We say that C' is *reachable in one step from* C , denoted $C \rightarrow C'$, if $C \xrightarrow{e} C'$ for some encounter $e \in E_I$. We say that C' is *reachable* from C and write $C \rightsquigarrow C'$, if there is a sequence of configurations $C = C_0, C_1, \dots, C_t = C'$, such that $C_i \rightarrow C_{i+1}$ for all $i, 0 \leq i < t$.

An *execution* is a finite or infinite sequence of configurations C_0, C_1, C_2, \dots , where C_0 is an initial configuration and $C_i \rightarrow C_{i+1}$, for all $i \geq 0$. A *fairness condition* is imposed on the adversary to ensure the protocol makes progress. An infinite execution is *fair* if for every pair of configurations C and C' such that $C \rightarrow C'$, if C occurs infinitely often in the execution then so does C' . In what follows, every execution of a NET will by definition considered to be fair.

We define the *output of a configuration* C as the graph $G(C) = (V, E)$ where $V = \{u \in V_I : C(u) \in Q_{out}\}$ and $E = \{uv : u, v \in V, u \neq v, \text{ and } C(uv) = 1\}$. In words, the output-graph of a configuration consists of those nodes that are in output states and those edges between them that are active, i.e. the active subgraph induced by the nodes that are in output states. The output of an execution C_0, C_1, \dots is said to *stabilize* (or *converge*) to a graph G if there exists some step $t \geq 0$ s.t. $G(C_i) = G$ for all $i \geq t$, i.e. from step t and onwards the output-graph remains unchanged. Every such configuration C_i , for $i \geq t$, is called *output-stable*. The *running time* (or *time to convergence*) of an execution is defined as the minimum such t (or ∞ if no such t exists). Throughout the chapter, whenever we study the running time of a NET, we assume that interactions are chosen by a *uniform random scheduler* which, in every

step, selects independently and uniformly at random one of the $|E_I| = n(n-1)/2$ possible interactions. In this case, the running time on a particular n and an initial set of active edges E becomes a random variable (abbreviated “r.v.”) $X_{n,E}$ and our goal is to obtain bounds on $\max_{n,E} \{e[X_{n,E}]\}$, where $e[X]$ is the *expectation* of the r.v. X . That is, the running time of a protocol is defined here as the maximum (also called *worst-case*) expected running time over all possible initial configurations. Note that the uniform random scheduler is fair with probability 1.

Definition 2. We say that an execution of a NET on n processes *constructs a graph* (or *network*) G , if its output stabilizes to a graph isomorphic to G .

Definition 3. We say that a NET \mathcal{A} *constructs a graph language* L with *useful space* $g(n) \leq n$, if $g(n)$ is the greatest function for which: (i) for all n , every execution of \mathcal{A} on n processes constructs a $G \in L$ of order at least $g(n)$ (provided that such a G exists) and, additionally, (ii) for all $G \in L$ there is an execution of \mathcal{A} on n processes, for some n satisfying $|V(G)| \geq g(n)$, that constructs G . Equivalently, we say that \mathcal{A} *constructs L with waste $n - g(n)$* .

In this work, we shall also be interested in NETs that construct a graph language and additionally always *terminate*.

Definition 4. We call a NET \mathcal{A} *terminating* (or say that \mathcal{A} *always terminates*) if every execution of \mathcal{A} reaches a *halting* configuration, that is one in which every node is in a state q_h from a set of halting states Q_{halt} , where $(q_h, q, s) \rightarrow (q_h, q, s)$ (i.e. is ineffective) for every $q_h \in Q_{halt}$, $q \in Q$, and $s \in \{0, 1\}$.

Finally, in order to consider TM simulations, we denote by $\mathbf{SSPACE}(f(n))$ the symmetric subclass of the complexity class $\mathbf{SPACE}(f(n))$.

3.2 Problem Definitions

Acyclicity. Let $G = (V, A)$ be the subgraph of G_I consisting of V and the active edges between nodes in V , that is $A = \{e \in E_I : C(e) = 1\}$. The initial G is connected. The goal is for the processes to stably transform G to an acyclic graph spanning V without ever breaking the connectivity of G .

Line Transformation. Let $G = (V, A)$ be the subgraph of G_I consisting of V and the active edges between nodes in V , that is $A = \{e \in E_I : C(e) = 1\}$. The initial G is connected. The goal is for the processes to stably transform G to a spanning line.

Terminating Line Transformation. The same as Line Transformation with the additional requirement that all processes must terminate.

3.3 Fundamental Inabilities

We now give a few basic impossibility results that justify the necessity of minimally strengthening the network constructors model in order to be able to solve the above main problems.

The following proposition (which is a well-known fact in the relevant literature but we include here a proof for self-containment) states that if the system does not involve edge states (i.e. the original population protocol model with transition function $\delta : Q \times Q \rightarrow Q \times Q$), then a protocol cannot decide with termination whether there is a single a in the population (mainly because a node does not know how much time it has to wait to meet every other node). Though the result is not directly applicable to our model, still we believe that it might help the reader's intuition w.r.t. to the computational difficulties in this family of models.

Proposition 1 (PPs Impossibility of Termination). *There is no population protocol that can compute with termination the predicate $(N_a \geq 1)$ (i.e. whether there exists an a in the input assignment).*

Proof. Consider a population of size n and let the nodes be u_1, u_2, \dots, u_n . It suffices to prove the impossibility for the variation in which there is a unique leader, initially in state l , and all other nodes are non-leaders, initially in state q_a if their input is a and q_b if their input is not a , and all interactions are between the leader and the non-leaders. This is w.l.o.g. because this model is not weaker than the original population protocol model, which means that an impossibility for this model also transfers to the original population protocol model. Indeed, this model can easily simulate the original model as follows. Interactions between two non-leaders can be simulated via the leader: the leader first collects the state q_1 of a node u , which it marks, and then waits to interact with another node v . When this occurs, if the state of v is q_2 , rule $(q_1, q_2) \rightarrow (q'_1, q'_2)$ is applied to the state stored by the leader and to the state of v . Then the leader waits to meet u again (which can be detected since it has been marked) in order to update its state to q'_1 . When this occurs, the leader drops the stored information and starts a new simulation round.

So, let u_1 be the initial leader. Let A be a protocol that computes $(N_a \geq 1)$ and terminates on every n and every input assignment. Consider now the input assignment in which all inputs are b , that is there is no a and thus all non-leaders begin with initial state q_b . Clearly, it must hold that in every fair execution the leader terminates in a finite number of steps and says “no”. These steps are interactions between the leader and the non-leaders so any such execution can be represented by a sequence of u_i s from $\{u_2, \dots, u_n\}$. Let now $s = v_1, v_2, \dots, v_k$ be any such finite execution in which the leader says “no”. $v_i \in \{u_2, \dots, u_n\}$ is simply the node with which the leader interacted at step i .

Consider now a population of size $n + 1$. The only difference to the previous setting is that now we have added a node u_{n+1} with input a . Since now the predicate evaluates to 1, in every fair execution, A should terminate in a finite number of steps and say “yes”. Take any fair execution $s' = s, v_{k+1}, \dots, v_h$, that is s' has s as an “un-fair” prefix. As s contains the same nodes as before with the same input assignment,

the leader in s' terminates in precisely k steps saying “no” without knowing that an additional node with input a exists in this case. This contradicts the existence of protocol A . We should mention that the leader has no means of guessing the existence of node u_{n+1} because its termination only depends on the protocol stored in its memory which is by definition finite and independent of n (suffices to consider the longest chain of rules that leads to termination with output “no” and which corresponds to at least one feasible execution). \square

Moreover, even if the system is initially connected, there are some very symmetric topologies that do not allow for strong computations. For example, if the topology is a star with the leader at the center, then the system is equivalent to population protocols on a complete interaction graph and can compute only semilinear predicates on input assignments, again only in an eventually-stabilizing way (i.e. no termination). This is captured by the following proposition.

Proposition 2 (Structure vs Computational Power). *There are initial topologies in which the computational power of population protocols is as limited as in the case of no structure at all.*

The above expose the necessity of additional assumptions, such as topology modifications, in order to hope for terminating computations and surpass the computational power of classical population protocols. So, we turn our attention again to our model, i.e. where the edges have binary states and the protocol can modify them, and consider the case in which the initial topology is always connected.

Proposition 3 (NETs Impossibility of Termination). *There is no protocol that can compute with termination the predicate $(N_a \geq 1)$, even if the initial topology is connected and even if there is a pre-elected unique leader.*

So, connectivity of the initial topology alone, even if the protocol is allowed to transform the topology, is not sufficient for non-trivial terminating computations. In the rest of the chapter we shall naturally try to overcome this by adding to the model minimal and realistic extra assumptions. Interestingly, it will turn out that there are some very plausible such assumptions that allow for: (i) termination and (ii) computation of all predicates on input assignments that can be computed by a TM in quadratic space ($O(n^2)$, where n is the number of nodes).

One of the assumptions that we will keep throughout is that the nodes are capable of detecting some small local degrees. For example, in Section 4 we will assume that a node can detect that it has local degree 1 or 2, otherwise it knows that it has degree in $\{0, 3, 4, \dots, n-1\}$ without being able to tell its precise value. We will complement this local degree detection mechanism with either a unique leader or a common neighbor detection mechanism in order to arrive at the above strong characterization.

Keep in mind that we want to give protocols for Acyclicity and Terminating Line Transformation. In Acyclicity, the protocol begins from any connected active topology and has to transform it to an acyclic network without ever breaking the connectivity, while in Terminating Line Transformation the protocol does not necessarily have to preserve connectivity but it has to satisfy the additional requirements its

constructed network to be a spanning line and to always terminate. Still, even for the Terminating Line Transformation problem we shall mostly focus on protocols that perform the transformation without ever breaking connectivity. A justification of this choice, is that arbitrary connectivity breaking could render the protocol unable to terminate even if the protocol is equipped with all the additional mechanisms mentioned above. This is made formal in Lemma 2 of Section 5.1. One way to appreciate this is to consider a protocol in which a leader breaks in some execution the network into an unbounded number of components. Then the leader can no longer distinguish an execution in which one of these components is being concealed from an execution that it is not. For example, if the leader is trying to construct a spanning line, then it has no means of distinguishing a spanning line on all nodes but the concealed ones from one on all nodes. Of course, this does not exclude protocols that perform some controlled connectivity breakings, e.g. a leader breaking a spanning line at one point and then waiting to reconnect the two parts. So, in principle, our problems could have been defined independently of whether connectivity is preserved or not, as they can also be solved in some cases by protocols that do not always preserve connectivity. However, in this work, for simplicity and clarity of presentation, we have chosen to focus only on those protocols that always preserve connectivity.

Before starting to present our protocols for above problems and the upper bounds on time provided by them, we give a lower bound on the time that any protocol needs in order to solve the Line Transformation problem.

Lemma 1 (Line Transformation Lower Bound). *The running time of any protocol that solves the Line Transformation problem is $\Omega(n^2 \log n)$.*

4 Transformers with a Unique Leader

We begin from the simplest case in which there is initially a pre-elected unique leader that handles the transformation. Recall that the initial active topology is connected. The goal is for the protocol to transform the active topology to a spanning line and when this occurs to detect it and terminate (i.e. solve the Terminating Line Transformation problem). Ideally, the transformation should preserve connectivity of the active topology during its course (or break connectivity in a controlled way, because, as we already discussed in the previous section, uncontrolled/arbitrary connectivity breaking may render termination impossible). Moreover, as a minimal additional assumption to make the problem solvable (in order to circumvent the impossibility of Proposition 3), we assume that a node can detect whether it has local degree 1 or 2 (otherwise it knows that it has degree in $\{0, 3, 4, \dots, n - 1\}$ without being able to tell its precise value). We first give a straightforward solution, with a complete presentation of its transitions and an illustration showing them in action. Though that protocol is correct, it is rather slow and it mainly serves as a demonstration of the model and the problem under consideration. Then we follow a different approach and arrive at a time-optimal protocol for the problem.

The idea of the first protocol is simple. The leader begins from its initial node and starts forming an arbitrary line by expanding one endpoint of the line towards unvisited nodes. Every such expansion either occurs over an edge that was already active from the very beginning or over an inactive edge which the protocol activates. Apart from expanding its active line with the goal of making it spanning after $n - 1$ expansions, the leader must also guarantee that eventually no cycles will have remained. One idea would be to first form a spanning line and then start eliminating all unnecessary cycles, however there is, in general, no way for the protocol to detect that the line is indeed spanning, due to the possible presence of non-line active edges joining nodes of the line. This is resolved by eliminating line-internal cycles “online” after every expansion of the line. This guarantees that when the last expansion occurs and the protocol deactivates the last cycles, the active topology will be a spanning line. Now the protocol can easily detect this by traversing the line from left to right and comparing the observed active degree sequence to the target degree sequence $1, 2, 2, \dots, 2, 1$ (i.e. the degree sequence of a spanning line). We next give the detailed description of the protocol.

Protocol Online-Cycle-Elimination. The leader marks its initial node as “left endpoint” e_l and picks an arbitrary next node for the line (for the first step it could be from its active neighbors, because there is at least one such node due to initial connectivity) and marks that node as “right endpoint” e_r . Then the leader moves to e_r , finds an arbitrary next node which is not part of the current line, if the edge is inactive it activates it and marks that node as e_r and the previous e_r is converted to i (for “internal node” of the line). Observe that the active line is always in special states, which makes its nodes detectable.

After every such expansion, the leader starts a cycle elimination phase. In particular, the leader deactivates all edges that introduce a cycle inside its active line. To do this, it suffices after every expansion to deactivate the cycles introduced by the new right endpoint e_r . First, the leader moves to e_l (e.g. by direct communication or by traversing the active line to the left). Every time, the leader waits to meet e_r , in order to check the status of the edge; if it is active, it deactivates it and then moves on step to the right on the line. When the leader arrives at the left neighbor of e_r , all line-internal cycles have been eliminated and the leader just moves to e_r . If the degree sequence observed during the traversal to the right (the degree of a node is checked *after* checking and possibly modifying the status of its edge to e_r) was of the form $1, 2, 2, \dots, 2, 1$ then the line is spanning and the leader terminates. Otherwise, the line is not spanning yet and the leader proceeds to the next expansion.

The code of the protocol is presented in Protocol 1. For readability, we only present the code for the expansion and cycle elimination phases and we have excluded the termination detection subroutine (it is straightforward to extend the code to also take this into account). An illustration showing what are the roles of the various states and transitions during the expansion and cycle elimination phases, is given in Figure 1.

Protocol 1 Online-Cycle-Elimination

$Q = \{l_0, l_1, l_c, l'_c, l, q_0, e, i, i', t, t_e, t_f, t'_f, t_r, t', p, p', p''\}$, initially the unique leader is in state l_0 and all other nodes are in state q_0
 δ :

$$\begin{array}{lll}
(l_0, q_0, 1) \rightarrow (e, l_1, 1) & (t_e, i, 1) \rightarrow (e, t, 1) & (p'', t', 1) \rightarrow (p, t, 1) \\
(l_1, q_0, \cdot) \rightarrow (i', l'_c, 1) & (t, l_c, \cdot) \rightarrow (t_r, l_c, 0) & (t_r, i', 1) \rightarrow (p', t'_f, 1) \\
(e, l'_c, \cdot) \rightarrow (t_e, l_c, 0) & (t_r, i, 1) \rightarrow (p', t', 1) & (p'', t'_f, 1) \rightarrow (i, t_f, 1) \\
(t_e, i', 1) \rightarrow (e, t_f, 1) & (e, p', 1) \rightarrow (e, p'', 1) & (l, q_0, \cdot) \rightarrow (i', l'_c, 1) \\
(t_f, l_c, 1) \rightarrow (i, l, 1) & (p, p', 1) \rightarrow (i, p'', 1) &
\end{array}$$

// All transitions that do not appear have no effect

// The logical structure is better followed if the transitions are read from top to bottom

Theorem 1. *By assuming a pre-elected unique leader and the ability to detect local degrees 1 and 2, Protocol Online-Cycle-Elimination solves the Terminating Line Transformation problem in $\Theta(n^4)$ time.*

Proof. We prove the following invariant: “For all $1 \leq i \leq n - 1$, after the i th expansion and cycle elimination phases, the leader lies on the e_r endpoint of an active line of (edge-)length i without line-internal cycles (still any node of the line may have active edges to the rest of the graph) and the active topology is connected”. This implies that for $i < n - 1$ there is at least one node of the line that has an edge to a node not belonging to the line and that for $i = n - 1$ the active topology is a spanning line (without any other active edges).

First observe that connectivity never breaks, because whenever the protocol deactivates an edge $e = uv$, both u and v are nodes belonging to the active line formed so far (in particular, at least one of them is the e_r endpoint of the line). As e is an edge forming a cycle on the active line after its deactivation connectivity between u and v still exists by traversing the line.

We prove by induction the rest of the invariant. It holds trivially for $i = 1$. Given that it holds for any $1 \leq i \leq n - 2$ we prove that it holds for $i + 1$. By hypothesis, when expansion $i + 1$ occurs, the only possible line-internal cycles are between the new e_r and the rest of the line. During the cycle elimination phase the protocol eliminates all these cycles, and as a result by the end of phase $i + 1$ the active line has now length $i + 1$, it has no internal cycles and is still connected to the rest of the graph.

It remains to show that the leader terminates just after phase $n - 1$ and never at a phase $i < n - 1$. For the first part, after phase $n - 1$ the active topology is a spanning line, thus the observed degree sequence when the leader traverses it from left to right is of the form $1, 2, 2, \dots, 2, 1$ which triggers termination. For the second part, after any phase $i < n - 1$ there is at least one node of the line having an active edge leading outside the line. In case that node is an endpoint, its active degree is at least 2 and in case it is an internal node its active degree is at least 3 (after eliminating a

possible cycle of that node with e_r). So, in this case the observed degree sequence is not of the form $1, 2, 2, \dots, 2, 1$ and, as required, the leader does not terminate.

For the running time, the worst case is when the initial active topology is a clique. In this case, the protocol must deactivate $\Theta(n^2)$ edges to transform the clique to a line. Every edge deactivation is performed by placing a mark on each endpoint of the edge and waiting for the scheduler to pick that edge for interaction. This takes time $\Theta(n^2)$, so the total time for deactivating $\Theta(n^2)$ edges is $\Theta(n^4)$. \square

We should mention that due to the unique-leader guarantee, it suffices to only have detection of whether the degree is equal to 1 (i.e. the detection of degree equal to 2 can be dropped). The reason is that the leader can every time break the line at some point while marking the two endpoints of the edge and then check whether one of these nodes has degree 1. If yes, then its previous degree was 2 and the leader waits for the two marked nodes to interact again in order to reconnect them, now knowing their degree.

A drawback of the above protocol is that it is rather slow. In the paper on which this chapter is based, we have developed another protocol, based on a different transformation technique, which is time-optimal. That protocol is called Line-Around-a-Star.

Theorem 2. *By assuming a pre-elected unique leader and the ability to detect local degree 1, Protocol Line-Around-a-Star solves the Terminating Line Transformation problem. Its running time is $\Theta(n^2 \log n)$, which is optimal.*

5 Transformers with Initially Identical Nodes

An immediate question, given the optimal Line-Around-a-Star protocol, is whether the unique leader assumption can be dropped and still have a correct and possibly also optimal protocol for Terminating Line Transformation. At a first sight it might seem plausible to expect that the problem is solvable. The reason is that the nodes can execute a leader election protocol (e.g. the standard pairwise elimination protocol; see e.g. [AR09]) guaranteeing that eventually a single leader will remain in the system which can from that point on handle the execution of one of the leader-based protocols of the previous section. The only additional guarantee is to ensure that nothing can go wrong as long as there are more than one leaders in the population. Typically, this is achieved in the population protocol literature by the reinitialization technique in which the configuration of the system is reinitialized/restored every time another leader is eliminated so that when the last leader remains a final reinitialization gives a correct system configuration for the leader to work on. In fact, this technique and others have been used in the population protocol literature to show that most population protocol models do not benefit in terms of computational power from the existence of a unique leader (still they are known to benefit in terms of efficiency).

In contrast to this intuition, we shall see in this section (see Corollary 2) that if all nodes are initially identical, Terminating Line Transformation becomes impossible to solve (with the modeling assumptions we have made so far). In particular, we will show that any protocol that makes the active topology acyclic, may disconnect it in some executions in $\Theta(n)$ components (see Corollary 1). As already discussed in Section 3.3, such a worst-case disconnection is severe for any terminating protocol, because, in this case, a component has no means of determining when it has interacted with (or heard from) all other components in the network.

Observation 1 *For a protocol to transform any topology to a line (or in general to an acyclic graph) without breaking connectivity, it must hold that the protocol deactivates an edge only if the edge is part of a cycle. Because deleting an edge e of an undirected graph does not disconnect the graph iff e is part of a cycle.*

There are several ways to achieve this when there is a unique leader. However, it will turn out that this is not the case when all nodes are initially identical.

5.1 Impossibility Results

An immediate question is whether there is a protocol with initially identical nodes that decides the existence of small cycles and additionally always terminates. We shall now show that this is not the case.

Theorem 3 (Strong Impossibility). *For every connected graph G with at least one cycle, there is an infinite family of graphs \mathcal{G} such that for every $G' \in \mathcal{G}$ every protocol (beginning from identical states on all nodes) that makes G acyclic may disconnect G' in some executions.*

The above strong result states that *every* connected graph G has a corresponding infinite family of graphs (in most cases disjoint to the families of other graphs) such that Acyclicity cannot be solved at the same time on G and on a G' from the family. This means that it does not just happen for Acyclicity to be unsolvable in a few specific inconvenient graphs. *All graphs* are in some sense inconvenient for Acyclicity when studied together with the families that we have defined.

Corollary 1 (Acyclicity Impossibility). *If all nodes are initially identical, then any protocol that always makes the active topology acyclic may disconnect it in some executions in $\Theta(n)$ active components (i.e. in a worst-case manner).*

Lemma 2. *If a protocol breaks in some executions the active topology into $\Theta(n)$ components, then such a protocol cannot solve the Terminating Line Transformation problem.*

Corollary 2 (Terminating Line Transformation Impossibility). *If all nodes are initially identical, there is no protocol for Terminating Line Transformation.*

5.2 *The Common Neighbor Detection Assumption*

In light of the impossibility results of the previous section, we naturally ask whether some minimal strengthening of the model could make the problems solvable. To this end, we give to the nodes the ability to detect whether they have a neighbor in common. In particular, we assume that whenever two nodes interact, they can tell whether they have at that time a common neighbor (over active edges). Clearly, this mechanism can be used to safely deactivate an edge in case it happens that the two nodes are indeed part of a 3-cycle. If the two nodes are part only of longer cycles they still cannot deactivate the edge with certainty. Observe that the common neighbor detection mechanism is very local and easily implementable by almost any plausible system. For example, it only requires local names and at least 2-round local communication before neighborhood changes. Moreover, it is also an inherent capability of the variation of population protocols in which the nodes interact in triples instead of pairs (see e.g. [AAD⁺06, BBRR12]). Interestingly, we shall see in this section that this minimal extra assumption overcomes the impossibility results both of Corollary 1 and Corollary 2. In particular, both Acyclicity and Terminating Line Transformation become now solvable.

Proposition 4. *By assuming that nodes are equipped with the common neighbor detection mechanism, there is a protocol, called Star-Transformer, that solves Acyclicity in the setting in which all nodes are initially identical. In particular, the final acyclic active topology is always a spanning star.*

We now exploit the common neighbor detection assumption and the Star-Transformer protocol to give a correct and efficient protocol for the Terminating Line Transformation problem. The protocol, called Line-Transformer, assumes (as did the protocols of Section 4) the ability to detect whether the local active degree of a node is equal to 1 or 2.

Protocol Line-Transformer. We give here a high-level description. All nodes are initially leaders in state l . When two leaders interact, one of them becomes a peripheral in state p and the edge is activated. Every leader is connected to all ps that it encounters. Two ps deactivate an active edge joining them only if at the time of interaction they have a neighbor in common. When a peripheral has active local degree equal to 1, its local state is p_1 , otherwise it is p or one of some other states that we will describe in the sequel.

When a leader first sees one of its own p_1s (i.e. via an active edge), it initiates the formation of a line over its p_1 peripherals (observe that the set of p_1 peripherals of a leader does not remain static, as e.g. a p_1 becomes p when some other leader connects to it). In particular, as in Protocol Line-Around-a-Star, the line will have as its “left” endpoint the center of the local star, which will be in a new state e_l , and it will start expanding over the available local p_1 peripherals over its right endpoint in state l' .

The new center e_l keeps connecting to new peripherals but it cannot become eliminated any more by other leaders. Pairwise eliminations only occur via any combi-

nation of l and l' . A local line expands over the local p_1 s as follows. When the right endpoint l' encounters a p_1 , which can occur only via an inactive edge, it expands on it only if the two nodes have a common neighbor (which can only be the center of the local star). If this is satisfied, the l' takes the place of the p_1 , leaving behind an i (for “internal node” of the line) and the edge becomes activated. Moreover, the center e_l deactivates every active edge it has with an i but not with the first peripheral of the line (so the first peripheral that the line uses must always be in a distinguished state i_1) and not with the l' right endpoint (because that edge is always needed for common neighbor detection during the next expansion).

Now, if the l' endpoint ever meets either another l' endpoint or an l center then one of them becomes deactivated. When it meets an l center we can always prefer to deactivate the l center because no line backtracking is required in this case. When an l center is deactivated, l simply becomes p and the edge becomes activated (it can never be a p_1 immediately but this is minor).

The most interesting case is when an l' loses from another l' . In this case, the eliminated l' becomes f . The role of f is to backtrack the whole local line construction by simply converting one after the other every i on its left to p and finally converting e_l again to p (again it cannot be a p_1 at the time of conversion). This backtracking process cannot fail because f has always a single i (or i_1) active neighbor, always the one on its left, while its right neighbor is no one initially and a p in all subsequent steps, so it knows which direction to follow. When the backtracking process ends, all the nodes of the local star are either p or p_1 so they can be attracted by the stars that are still alive.

The protocol terminates, when for the first time it holds that an e_l has local degree equal to 2 after its line has for the first time length at least 3 (nodes). When this occurs, a spanning ring has been formed and e_l can deactivate the edge (e_l, l') between the two endpoints to make it a spanning line. This completes the description of the protocol.

Theorem 4. *By assuming that nodes are equipped with a common neighbor detection mechanism and have the ability to detect local degrees 1 and 2, Protocol Line-Transformer solves the Terminating Line Transformation problem in the setting in which all nodes are initially identical. Its running time is $O(n^3)$.*

Table 1 summarizes all protocols that we developed for the Terminating Line Transformation problem, both for the case of a pre-elected unique leader (Protocols Online-Cycle-Elimination and Line-Around-a-Star in Section 4) and for the case of identical nodes (Protocol Line-Transformer in the present section).

Finally, in the paper on which this chapter is based, we have shown how the spanning line formed with termination by *Line-Transformer* can be used to establish that the class of computable predicates is the maximum that one can hope for in this family of models.

Theorem 5 (Full Computational Power). *Let the initial active topology be connected, all nodes be initially identical, and let the nodes be equipped with degree in $\{1, 2\}$ detection and common neighbor detection. Then for every predicate $p \in \text{SSPACE}(n^2)$ there is a terminating NET that computes p .*

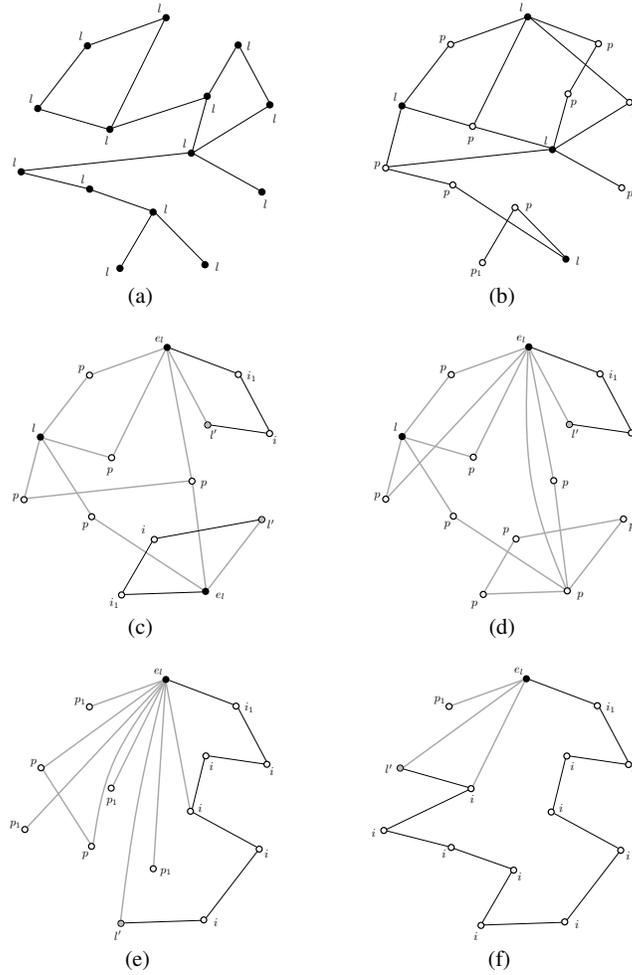


Fig. 2 An example execution of Protocol Line-Transformer. In all subfigures, black and gray edges are active and missing edges are inactive. Black and gray are used together whenever we want to highlight some subnetwork of the active network. (a) Initially, all nodes are leaders and the topology is connected. (b) Most leaders have been converted to peripherals, some leaders have attracted new peripherals, and some peripherals have disconnected from each other. (c) Two of the survived leaders have started to form lines over their p_1 peripherals. The centers of these stars are now in state e_l (black nodes), the other endpoint of their lines is in state l' (gray nodes), and the lines are drawn by black edges. (d) The l' endpoints of the two previous lines interacted and one of them was backtracked. (e) A single line has remained. (f) The line is almost spanning.

Protocol	Leader	DD	CND	Expected Time	Lower Bound
OCE	Yes	1	No	$\Theta(n^4)$	$\Omega(n^2 \log n)$
LAS	Yes	1	No	$\Theta(n^2 \log n)$ (opt)	$\Omega(n^2 \log n)$
LT	No	1,2	Yes	$O(n^3)$	$\Omega(n^2 \log n)$

Table 1 All protocols developed in this work for the Terminating Line Transformation problem. For each of these protocols (OCE: Online-Cycle-Elimination, LAS: Line-Around-a-Star, and LT: Line-Transformer), the table shows whether it makes use of a pre-elected unique leader, what local degree detection it uses (DD), whether it uses common neighbor detection (CND), and also its expected running time under the uniform random scheduler. The last column shows the best known lower bound for the problem.

6 Conclusions and Further Research

There are many open problems related to the findings of the present work. We have shown that initial connectivity of the active topology combined with the ability of the protocol to transform the topology yield, under some additional minimal and local assumptions, an extremely powerful model. We managed to show this by developing protocols that transform the initial topology to a convenient one (in our case the spanning line) while always *preserving the connectivity of the topology*. Though arbitrary connectivity breaking makes termination impossible, still we have not excluded the possibility that some protocol performs some “controlled” connectivity breaking during its course, being always able to correctly reassemble the disconnected parts and terminate.

Another issue has to do with the underlying interaction model. Throughout this work we have assumed that the underlying interaction graph is the clique K_n and all of our protocols largely exploit this. Though this model is a convenient starting point to understand the basic principles of algorithmic transformations of networks, it is obvious that it is highly non-local. Realistic implementations would probably require more local or geometrically constrained models (like the one of [Mic15]), for example, one in which, at any given time, a node can only communicate with nodes at active distance at most 2. It is also valuable to consider the Terminating Line Transformation and Acyclicity problems in models of computationally weak (and probably also anonymous) robots moving in the plane.

There are also some more technical intriguing open questions. The most prominent one is whether protocol Line-Transformer is time-optimal. Recall that its running time was shown to be $O(n^3)$. First of all, it is not clear whether the analysis is tight. The subroutine that dominates the running time is the one that tries to form a spanning line over the peripheral nodes, which is restricted by the fact that the partial lines of “sleeping” stars have to either be backtracked (which is what our solution does) or merged somehow with the lines of “awake” stars. We should mention that the spanning line subroutine that backtracks many “sleeping” lines in parallel is an immediate improvement of the best spanning line protocol of [MS14], called Fast-Global-Line. The improvement is due to the fact that instead of having the

awake leader backtrack node-by-node sleeping lines, we now have any sleeping line backtrack itself, so that many backtrackings occur in parallel. We also have experimental evidence showing a small improvement [ALMS15] but still we do not have a proof of whether this is also an asymptotic improvement. For example, is it the case that the running time of this improvement is $O(n^3/\log n)$ (or even smaller)? This question is open. There is also room for lower bounds. Apart from the obvious lower bound for the Terminating Line Transformation problem with identical nodes, one could also focus on the spanning line construction problem with initially disconnected nodes (i.e. the Spanning Line problem of [MS14]). The reason is that an improvement to this problem would probably imply an improvement for Terminating Line Transformation by using the protocol as an improved subroutine of Line-Transformer for forming the lines over the peripherals of the star. The best lower bound known for Spanning Line is $\Omega(n^2)$. Some first attempts suggest that it might be non-trivial to improve this to $\Omega(n^2 \log n)$.

References

- AAC⁺05. D. Angluin, J. Aspnes, M. Chan, M. J. Fischer, H. Jiang, and R. Peralta. Stably computable properties of network graphs. In *1st IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS)*, volume 3560 of LNCS, pages 63–74. Springer-Verlag, June 2005.
- AAAD⁺06. D. Angluin, J. Aspnes, Z. Diamadi, M. J. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18[4]:235–253, March 2006.
- AAER07. D. Angluin, J. Aspnes, D. Eisenstat, and E. Ruppert. The computational power of population protocols. *Distributed Computing*, 20[4]:279–304, November 2007.
- ACD⁺11. G. Aloupis, S. Collette, M. Damian, E. D. Demaine, R. Flatland, S. Langerman, J. O’Rourke, V. Pinciu, S. Ramaswami, V. Sacristán, and S. Wuhrer. Efficient constant-velocity reconfiguration of crystalline robots. *Robotica*, 29[01]:59–71, 2011.
- ALMS15. D. Amaxilatis, M. Logaras, O. Michail, and P. G. Spirakis. NETCS: A new simulator of population protocols and network constructors. *arXiv preprint arXiv:1508.06731*, 2015.
- AR09. J. Aspnes and E. Ruppert. An introduction to population protocols. In B. Garbinato, H. Miranda, and L. Rodrigues, editors, *Middleware for Network Eccentric and Mobile Applications*, pages 97–120. Springer-Verlag, 2009.
- BBRR12. J. Beauquier, J. Burman, L. Rosaz, and B. Rozoy. Non-deterministic population protocols. In *16th International Conference on Principles of Distributed Systems (OPODIS)*, LNCS, pages 61–75. Springer, 2012.
- CDF⁺09. I. Chatzigiannakis, S. Dolev, S. P. Fekete, O. Michail, and P. G. Spirakis. Not all fair probabilistic schedulers are equivalent. In *13th International Conference on Principles of Distributed Systems (OPODIS)*, volume 5923 of *Lecture Notes in Computer Science*, pages 33–47. Springer-Verlag, 2009.
- CKLWL09. A. Cornejo, F. Kuhn, R. Ley-Wild, and N. Lynch. Keeping mobile robot swarms connected. In *Proceedings of the 23rd International Symposium on Distributed Computing (DISC)*, LNCS, pages 496–511. Springer, 2009.
- CMN⁺11. I. Chatzigiannakis, O. Michail, S. Nikolaou, A. Pavlogiannis, and P. G. Spirakis. Passively mobile communicating machines that use restricted space. *Theoretical Computer Science*, 412[46]:6469–6483, October 2011.

- CMNS13. I. Chatzigiannakis, O. Michail, S. Nikolaou, and P. G. Spirakis. The computational power of simple protocols for self-awareness on graphs. *Theoretical Computer Science*, 512:98–118, November 2013.
- DFSY15. S. Das, P. Flocchini, N. Santoro, and M. Yamashita. Forming sequences of geometric patterns with oblivious mobile robots. *Distributed Computing*, 28[2]:131–145, April 2015.
- Dot14. D. Doty. Timing in chemical reaction networks. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 772–784, 2014.
- MCS11a. O. Michail, I. Chatzigiannakis, and P. G. Spirakis. Mediated population protocols. *Theoretical Computer Science*, 412[22]:2434–2450, May 2011.
- MCS11b. O. Michail, I. Chatzigiannakis, and P. G. Spirakis. *New Models for Population Protocols*. N. A. Lynch (Ed), Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool, 2011.
- Mic15. O. Michail. Terminating distributed construction of shapes and patterns in a fair solution of automata. In *Proceedings of the 34th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 37–46. ACM, 2015.
- MS14. O. Michail and P. G. Spirakis. Simple and efficient local codes for distributed stable network construction. In *Proceedings of the 33rd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 76–85. ACM, 2014. Also in *Distributed Computing*, doi: 10.1007/s00446-015-0257-4, 2015.
- MS15. O. Michail and P. G. Spirakis. Terminating population protocols via some minimal global knowledge assumptions. *Journal of Parallel and Distributed Computing (JPDC)*, 81:1–10, 2015.
- SY99. I. Suzuki and M. Yamashita. Distributed anonymous mobile robots: Formation of geometric patterns. *SIAM J. Comput.*, 28[4]:1347–1363, March 1999.
- WCG⁺13. D. Woods, H.-L. Chen, S. Goodfriend, N. Dabby, E. Winfree, and P. Yin. Active self-assembly of algorithmic shapes and patterns in polylogarithmic time. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 353–354. ACM, 2013.